



Calhoun: The NPS Institutional Archive
DSpace Repository

Theses and Dissertations

1. Thesis and Dissertation Collection, all items

2014-06

Developing standard exercises and statistics to measure the impact of cyber defenses

Berninger, Matthew L.

Monterey, California: Naval Postgraduate School

<http://hdl.handle.net/10945/42581>

Copyright is reserved by the copyright owner.

Downloaded from NPS Archive: Calhoun



Calhoun is the Naval Postgraduate School's public access digital repository for research materials and institutional publications created by the NPS community. Calhoun is named for Professor of Mathematics Guy K. Calhoun, NPS's first appointed -- and published -- scholarly author.

Dudley Knox Library / Naval Postgraduate School
411 Dyer Road / 1 University Circle
Monterey, California USA 93943

<http://www.nps.edu/library>



NAVAL POSTGRADUATE SCHOOL

MONTEREY, CALIFORNIA

THESIS

**DEVELOPING STANDARD EXERCISES AND
STATISTICS TO MEASURE THE IMPACT OF CYBER
DEFENSES**

by

Matthew L. Berninger

June 2014

Thesis Advisor:
Co-Advisor:

John Krautheim
Garrett McGrath

Approved for public release; distribution is unlimited

THIS PAGE INTENTIONALLY LEFT BLANK

REPORT DOCUMENTATION PAGE			<i>Form Approved OMB No. 0704-0188</i>	
Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instruction, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188) Washington, DC 20503.				
1. AGENCY USE ONLY (Leave blank)		2. REPORT DATE June 2014	3. REPORT TYPE AND DATES COVERED Master's Thesis	
4. TITLE AND SUBTITLE DEVELOPING STANDARD EXERCISES AND STATISTICS TO MEASURE THE IMPACT OF CYBER DEFENSES			5. FUNDING NUMBERS	
6. AUTHOR(S) Matthew L. Berninger				
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Naval Postgraduate School Monterey, CA 93943-5000			8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING /MONITORING AGENCY NAME(S) AND ADDRESS(ES)			10. SPONSORING/MONITORING AGENCY REPORT NUMBER	
11. SUPPLEMENTARY NOTES The views expressed in this thesis are those of the author and do not reflect the official policy or position of the Department of Defense or the U.S. Government. IRB Protocol number ____N/A____.				
12a. DISTRIBUTION / AVAILABILITY STATEMENT Approved for public release; distribution is unlimited			12b. DISTRIBUTION CODE	
13. ABSTRACT (maximum 200 words) <p>As companies seek protection from cyber attacks, justifying proper levels of investment in cyber security is essential. Like all investments, cyber defense costs must be weighed against their expected benefits.</p> <p>While some cyber investment models exist that can relate costs and benefits, these models are largely untested with experimental data. This research develops an experimental framework and statistics for testing and measuring the efficacy of cyber mitigation methods, such that they can be integrated into existing cyber investment models.</p> <p>This work surveys cyber security investment models and frameworks. Using cyber exercises as a source of attack data, types of exercises and how information is recorded was studied.</p> <p>A proof of concept for an experimental framework able to record statistics on cyber exercise attacks and defenses was developed. The environment is intended to resemble that of an actual cyber attack, and to collect attack and defense data in a repeatable and technology-agnostic manner. Possible future work could illuminate mathematical relationships between threat and mitigation.</p> <p>Statistics and procedures are proposed that are applicable to the specific proposed and similar frameworks. Such statistics could be incorporated into cyber models, ultimately leading to a more rational understanding of cyber attack and defense.</p>				
14. SUBJECT TERMS Incident Response, Cyber Exercises, Metrics, Cyber Investment Modeling			15. NUMBER OF PAGES 77	
			16. PRICE CODE	
17. SECURITY CLASSIFICATION OF REPORT Unclassified	18. SECURITY CLASSIFICATION OF THIS PAGE Unclassified	19. SECURITY CLASSIFICATION OF ABSTRACT Unclassified	20. LIMITATION OF ABSTRACT UU	

THIS PAGE INTENTIONALLY LEFT BLANK

Approved for public release; distribution is unlimited

**DEVELOPING STANDARD EXERCISES AND STATISTICS TO MEASURE
THE IMPACT OF CYBER DEFENSES**

Matthew L. Berninger
Civilian, Department Of Homeland Security
B.A., Columbia University, 2008

Submitted in partial fulfillment of the
requirements for the degree of

MASTER OF SCIENCE IN CYBER SYSTEMS AND OPERATIONS

from the

**NAVAL POSTGRADUATE SCHOOL
June 2014**

Author: Matthew L. Berninger

Approved by: John Krautheim
Thesis Advisor

Garrett McGrath
Co-Advisor

Cynthia Irvine, Ph.D.
Chair, Department of Cyber Academic Group

THIS PAGE INTENTIONALLY LEFT BLANK

ABSTRACT

As companies seek protection from cyber attacks, justifying proper levels of investment in cyber security is essential. Like all investments, cyber defense costs must be weighed against their expected benefits.

While some cyber investment models exist that can relate costs and benefits, these models are largely untested with experimental data. This research develops an experimental framework and statistics for testing and measuring the efficacy of cyber mitigation methods, such that they can be integrated into existing cyber investment models.

This work surveys cyber security investment models and frameworks. Using cyber exercises as a source of attack data, types of exercises and how information is recorded was studied.

A proof of concept for an experimental framework able to record statistics on cyber exercise attacks and defenses was developed. The environment is intended to resemble that of an actual cyber attack, and to collect attack and defense data in a repeatable and technology-agnostic manner. Possible future work could illuminate mathematical relationships between threat and mitigation.

Statistics and procedures are proposed that are applicable to the specific proposed and similar frameworks. Such statistics could be incorporated into cyber models, ultimately leading to a more rational understanding of cyber attack and defense.

THIS PAGE INTENTIONALLY LEFT BLANK

TABLE OF CONTENTS

I.	INTRODUCTION.....	1
A.	BACKGROUND	1
B.	MOTIVATION	1
C.	PROBLEM STATEMENT	4
D.	APPROACH.....	5
E.	ASSUMPTIONS AND LIMITATIONS	5
F.	THESIS ORGANIZATION.....	6
II.	BACKGROUND—EXISTING CYBER MODELS	7
A.	THE COST OF FAILURE.....	7
1.	Stock Market Returns	7
2.	Cyber Insurance.....	8
B.	MEASURING BENEFITS.....	9
1.	The Gordon-Loeb Model.....	9
2.	Empirical Analysis	10
C.	WHERE IS THE DATA?.....	12
1.	Existing Metrics	12
2.	Exercises Hold Untapped Data	13
III.	EXERCISE METHODOLOGY AND DESIGN CONSIDERATIONS.....	15
A.	CYBER ATTACK METHODOLOGY	15
B.	CYBER EXERCISE FRAMEWORKS	17
C.	CONDITIONS OF THE PROPOSED EXERCISE FRAMEWORK.....	20
D.	STATISTICAL MEASURES	23
1.	Indicators and Infrastructure.....	23
2.	Tool-Based Statistics.....	24
E.	INCORPORATING STATISTICS INTO CYBER INVESTMENT MODEL	24
1.	Evaluating Alerting Technologies – Theory	24
2.	Properties of Indicators.....	26
3.	Implementation Into a Simplified Gordon-Loeb Model	26
4.	Detection Impact	27
5.	Summary.....	29
IV.	METHODOLOGY AND DATA GATHERING PROOF OF CONCEPT	31
A.	FRAMEWORK FOR ONE OR TWO-PLAYER EXERCISE.....	31
B.	PROCESS/ RESULTS.....	34
C.	MONITORING/ GATHERING STATISTICS	39
D.	EXTENSIONS TO CLASSROOM EXERCISE.....	42
V.	CONCLUSIONS AND FUTURE WORK.....	47
A.	LESSONS LEARNED	47
B.	FUTURE WORK.....	47
C.	CONCLUSIONS	50

LIST OF REFERENCES	53
INITIAL DISTRIBUTION LIST	57

LIST OF FIGURES

Figure 1.	Defcon-Style Capture-the-Flag, (from Cowan, Arnold, Beattie, Wright, & Viega, 2003).....	18
Figure 2.	Single or Dual User Exercise Topology	32
Figure 3.	Mapping the Network	34
Figure 4.	Sessions Started on Web Server	35
Figure 5.	Uploading the Exploit PDF.....	36
Figure 6.	Mapping Network Drive	38
Figure 7.	View/ Exfiltrate Sensitive Information.....	38
Figure 8.	IDS Alerts	39
Figure 9.	Avast! AV Alerts	41
Figure 10.	Classroom Exercise Topology	43
Figure 11.	Mitigation Timeline	50

THIS PAGE INTENTIONALLY LEFT BLANK

LIST OF TABLES

Table 1.	Hacking Exposed “Anatomy of a Hack” (from Scambray, McClure, & Kurtz, 2012)	15
Table 2.	Mandiant APT1 Methodology	16
Table 3.	Phases of Offense.....	22
Table 4.	Team-Based Measures and Statistics.....	23
Table 5.	Tool-Based Measures and Statistics	24
Table 6.	Calculation of Detection Impact	28
Table 7.	Calculation of Weighting Factor.....	29
Table 8.	Calculation of IDS Detection Impact.....	40
Table 9.	Calculation of Avast Detection Impact.....	41
Table 10.	Classroom Exercise Phases.....	45
Table 11.	Time-Based Statistics.....	48
Table 12.	Mitigation Statistics	49

THIS PAGE INTENTIONALLY LEFT BLANK

LIST OF EQUATIONS

Equation 1.	Simplified Gordon-Loeb	26
Equation 2.	Detection Impact	27
Equation 3.	Indicator Signal To Noise Ratio	27
Equation 4	Indicator Weight Factor	28
Equation 5	Breakout of Detection Impact	28
Equation 6.	Mitigation Timeliness	49

THIS PAGE INTENTIONALLY LEFT BLANK

LIST OF ACRONYMS AND ABBREVIATIONS

APT	advanced persistent threat
ARP	Address Resolution Protocol
AV	Anti-Virus
CIP	critical infrastructure protection / common industrial protocol
DHS	Department of Homeland Security
GUI	graphical user interface
IP	Internet protocol
IPS	intrusion prevention system
IT	information technology
NIST	National Institute of Standards and Technology
PCAP	packet capture
RDP	remote desktop protocol
TTP	tactics, techniques, and procedures

THIS PAGE INTENTIONALLY LEFT BLANK

ACKNOWLEDGMENTS

My sincere gratitude goes to Dr. Irvine, Dr. Davis, and the entire Cyber Academic Group at the Naval Postgraduate School for their guidance. Special thanks also go to John Krautheim, Scott Cote, and Garrett McGrath for my education in Red/Blue team exercises. Additionally, special thanks to Jay Holcomb and Davis Hake for my introduction to hacking demos and test environments. Thanks to my colleagues Nicholas Carr, Sean Spaniol, Mike King, Christopher Hallenbeck, Eric Fildebrandt, Ken Melton, and the rest of the NCCIC/US-CERT/ICS-CERT teams. Lastly, thanks to my family, friends, and lovely girlfriend for listening to me blabber on about cyber modeling and statistics.

THIS PAGE INTENTIONALLY LEFT BLANK

I. INTRODUCTION

A. BACKGROUND

Many senior officials in the U.S. Government and private sector contend that the sustained, strategic, and surreptitious theft of America's economic innovation could currently pose an existential threat to the economy and security of the United States (Rogers, 2011). As put by General Keith Alexander, the aggregate theft of American Intellectual property constitutes the "greatest transfer of wealth in history" (Rogin, 2012). While physical property is relatively secure, and the security of financial data has been made paramount, the intellectual property that fuels America's economy remains highly vulnerable. In order to understand this current state of cyber defense, researchers must examine the decisions that network defenders make, and the information necessary to make those decisions in an informed manner.

In order for senior leadership to make informed decisions, cyber defenders sorely need quantitative models and statistics. Real-world intrusion data is not readily available, nor will it become readily available in the near term. However, cyber defense exercises are free of much of the legal and classification restrictions of real intrusions, and can provide valuable data in a free, public, and standardized manner. The offensive and defensive tactics used during cyber exercises also mirror some of the procedures used in real world attacks. Therefore, exercises can serve as a risk-free proxy for data gathering; that data can then be used to feed cyber defense models, inform senior leadership decision-making, and further the understanding of the cybersecurity ecosystem.

B. MOTIVATION

In "Why Information Security is Hard," Ross Anderson outlines some of the possible reasons for the current "state" of cybersecurity defense. He argues that the motivations, liabilities, and responsibilities of parties in cyberspace largely determine which actions are taken, and by whom, to secure their networks. He points to "The

Tragedy of the Commons” as an allegorical example of the issue of collective action (or lack thereof) in cybersecurity:

If a hundred peasants graze their sheep on the village common, then whenever another sheep is added its owner gets almost the full benefit, while the other ninety-nine suffer only a small decline in the quality of the grazing. So they aren’t motivated to object, but rather to add another sheep of their own and get as much of the grazing as they can. The result is a dustbowl; and the solution is regulatory rather than technical. (Anderson, 2001)

The Tragedy of the Commons highlights the aggregated negative externalities of single actors’ self-interested actions. In cybersecurity, the issue is the same, but inverted: the lack of action on the part of a small fraction of companies to secure their networks makes everyone else more vulnerable. Citing Hal Varian’s 2000 NYTimes article, “Managing Online Security Risks” (Varian, 2000), Anderson argues that costs should be pushed to those “hosting” the malicious activity, so to speak. In the NYTimes, Varian puts it simply: “One reason that computer security is so poor in practice is that the liability is so diffuse.”

Adding to the discussion of forces at play in cyberspace, Andrew Oldyzko examines the “Economics, Psychology, and Sociology of Security” in his paper by the same name (Oldyzko, 2003). He touches on some of the economic motivations cited by Anderson, writing “security does not come for free, and so it is necessary to look at the tradeoffs between costs and benefits. Furthermore, it is necessary to look at the incentives of various players, as many have an interest in passing on the costs of security to others, or of using security for purposes such as protecting monopolies.” This highlights the importance of organizational and human interests behind decisions in the cyber defense realm. Furthermore, he believes the formal structures are, on paper, incomplete, writing: “Our commercial, government, and academic enterprises are large organizations with many formal rules and regulations. Yet the essential workings of these enterprises are typically based on various social relations and unwritten rules.” While technology and systems may have certain expected behaviors in a vacuum, the real performance of these technologies is at the mercy of the humans operating them, and by extension, their own cost-benefit decisions.

Information Sharing

In “Cyber Information-Sharing Models: An Overview” (The MITRE Corporation, 2012), researchers at MITRE argue that the current state of affairs in information sharing is a contributing factor to the poor state of cyber defense. They write that “a key element in defending against [cyber] attacks is having information about the tools, techniques and resources (physical, financial, and human) that adversaries are using to breach cyber defenses.” The argument here is that even with the right motivations, defense technologies, and security procedures in place, many organizations may simply lack the information they need to defend themselves.

These examples illustrate that there is much more to the cybersecurity problem than just poor cryptography or out-of-date AV signatures. There are large-scale economic, sociological, and information dissemination/translation issues at play, and these issues may not be addressed singly through technologies, policies, or legislation. Rather, it is critical to first understand the existing frameworks in which cybersecurity is currently discussed.

Current Paradigms

Though no domain or criminal area is quite like cyberspace, a discussion of cybersecurity models can be framed by existing paradigms. Perhaps the first macroeconomic paradigm that comes to mind is that of physical-world crime. Criminal economics has been studied for centuries, and many models exist, from narcotics to bank robbery to auto theft. If focused on cybercrime, these models may yield useful frameworks and bring researchers closer to understanding the mechanics and strategies of the cyber cartels. However, the existing criminal models are still largely rooted in the physical world, and thus may not account for the distributed and instantaneous nature of cyberspace.

Second is a military offense/defense model. Counting “good guys” on one side and “bad guys” on the other may help us define who is currently “winning.” This model,

too, is rooted in geography and may not correctly weigh the impact of small-scale issues, instead focusing on high-profile events and advanced capabilities.

A third paradigm to model cybersecurity is that of public health. Medical based models, specifically for infectious disease, exhibit the speed and networked environment of cyberspace. They also take into account the effect of mitigation measures on a population, including possible second-and third level after effects. Additionally, public health model demonstrates the interchangeable nature of threats and mitigations. That is, a given mitigation can be used to treat or deter a variety of threats. Finally, public health also shares the distributed and simultaneous nature of cyber defense—a given threat can strike a number of victims all around the same time, requiring rapid response and information sharing.

C. PROBLEM STATEMENT

In order to make a rational argument for cybersecurity investment, organizations must first ask whether cyber breaches present a *real* cost to the organization. Next, they may ask if insuring against those losses is a better solution versus actually trying to prevent them in the first place. Finally, if they do believe that the costs are worth preventing, and that insurance does not present a viable option, they must ask *which* measures to take, and what the expected benefits of those measures are. In order to make this assessment objectively, firms need to be able to perform a cost-benefit analysis of cyber security investments and measures. To do so, they must have an understanding of the expected benefits of cyber mitigation measures, as well as the expected cost of not applying those measures. Without this data, firms may not be able to empirically justify proactive cybersecurity investments.

However, this data is hard to find. Even if available, it is often not standardized. And lastly, it usually applies to a more macro understanding of the incident, rather than the specific tactics used by both offense and defense during an intrusion. While these details may not concern senior leadership, they do form the foundation of understanding cyber attacks and mitigations. That understanding can then inform strategic models and

decisions. Without this understanding, however, leaders are left to a theoretical expectation, or best guess estimate for their investments.

In kinetic warfare and physical crime, real-world simulations are understandably tricky. You cannot fully use real weapons in a wargame, and you cannot fire real bullets in a police exercise. In cyber, however, you can - and you can do it over and over again. For example, it is possible to build an environment which closely resembles a business network, complete with Windows domains, a DMZ network, routers, firewalls and other elements. This network can be attacked, infiltrated and destroyed without any real harm to real-world networks or operations. Using virtual machines, it can be rebuilt and done all over again. Cyber exercises, which already exist in many forms, thus present an untapped resource of cyber attack and defense statistics; we simply have not been collecting them properly. If efforts can be made to capture and standardize statistics during cyber exercises, the aggregate dataset could be used to build and test and model a myriad of scenarios.

D. APPROACH

This thesis project involves the development of a cyber exercise methodology in which a single user, two users, or classroom could create, gather, and standardize cyber attack and defense statistics. The objective is to develop a standardized exercise framework to gather statistics in a technology-agnostic manner. The exercise will be used to inform and test the statistics. While in theory, some statistics may make logical sense, if they cannot actually be collected and analyzed from an exercise, then they will not be usable in cyber investment calculations.

E. ASSUMPTIONS AND LIMITATIONS

This thesis assumes certain understanding of cyber security terms and concepts, as well as concepts used in mathematics and statistics. It is limited mainly by budget, hence the use of mainly open source and unsupported software in the creation of the attack lab.

The thesis project will be comprised of the discussion of cyber metrics and models, the reasons for using cyber exercise data to test those models, and finally the creation and testing of an experimental cyber attack framework.

While knowledge of cyber attack and defense methodologies is informed by real operational experience from US-CERT cases, all the data in this thesis is unclassified and fabricated in a lab environment. The tools and exploits used herein are seen and used in real-world environments.

F. THESIS ORGANIZATION

Chapter II provides background on existing cyber models and metrics. Chapter III provides the criteria for qualifying cyber exercises as a source of real statistical data. Chapter IV outlines the construction of the exercise environment, execution of a test exercise, and the results of data collection. Chapter V summarizes findings and proposes future work.

II. BACKGROUND—EXISTING CYBER MODELS

A. THE COST OF FAILURE

In a recently published Forrester Research Paper, “Determine The Business Value of An Effective Security Program —Information Security Economics 101,” Ed Ferrara argues that “To fully understand security’s financial impact on the organization, CISOs should understand all the various costs of protecting information. This includes the fixed costs as well as variable costs, especially those related to breaches” (Ferrara, 2002).

1. Stock Market Returns

One rather obvious—and elegant—method to measure the cost of a cybersecurity breach in the private sector is, simply, to measure the stock market impact of breach announcements. This is precisely what Katherine Campbell et al. studied in 2003. (Katherine Campbell, 2003) In the study, the researchers found “an overall negative stock market reaction to public announcements of information security breaches.” This is largely unsurprising news; however, Campbell further note, “We find a highly significant negative market reaction for information security breaches involving unauthorized access to confidential data, but no significant reaction when the breach does not involve confidential information.” Thus, the nature of the breach, in the eyes of investors, does appear to affect the market cost.

An associated study was performed by Gordon, Loeb, and Zhou in 2011 (Lawrence A. Gordon, 2011). The authors note conflicting studies since the 2003 study, saying, “several researchers have studied the impact of [cyber] breaches on the stock market returns of firms. The results from these studies, however, have been mixed.” Thus, Gordon et al. used a more “sophisticated market model over a long period and over two distinct and naturally arising sub-periods,” and in this study, attempted to “resolve conflicting evidence from previous studies concerning the effect of information security breaches.” (Lawrence A. Gordon, 2011). Their findings supported their earlier thoughts, that “...news about information security breaches (where such breaches are treated as a generic category) had a statistically significant effect on the stock returns of publicly

traded firms.” This appears to suggest, at least, that cyber breaches, in fact, can cause a very real negative financial impact to firms.

The 2011 stock-return analysis above also noted an interesting phenomenon: an apparent decrease in the market effect of a breach—namely, that breaches before 9/11 cost more to the stock market performance of a firm than breaches after 9/11 did. Since the breaches themselves had little other differences, human perception likely has much to do with this. Put by Gordon et al., “there seems to a shift in the attitude among investors toward viewing information security breaches as creating a corporate “nuisance” (or merely another recurring operating cost) rather than creating a potentially serious economic threat to the survival of firms.” If cyber breaches are, indeed, viewed as an “operating cost” in today’s environment, that may give private sector leaders some relief—after all, negative events can be covered by insuring against them.

2. Cyber Insurance

Despite its seemingly logical place in the cybersecurity field, cyber insurance has so far failed to provide a mature, widely accepted framework for insuring against cyber risk. In “Aegis, A Novel Cyber Insurance Model,” for example, Pal, Golubchik and Psounis propose a new quantitative model for cyber insurance, but also concede that their model, the Aegis framework,”also faces problems [related to] interdependent security.” (Pal, Golubchik, & Psounis, 2011). They are essentially saying that the fundamental interdependence of assets in cyberspace, among other things, currently limits the application of their cyber insurance model. That is not to say that the problem is intractable: Rainer Bohme and Gaurav Kataria attempt to model correlation in “Models and Measures for Correlation in Cyber-Insurance.” In this study, they find that because of the high degree of interdependence present in cyberspace, “our simulation results indicate that cyber-insurance is best suited for classes of risk with high internal and low global correlation,” that is, companies or assets which have minimal externalities. Additionally, they cite a lack of good empirical data with which to test their findings, suggesting that as a direction for future work. (Kataria, 2006)

In their 2010 paper, “Modeling Cyber-Insurance: Towards A Unifying Framework” (Bohme & Schwartz, 2010), Rainer, Bohm, and Schwartz examine various cyber-insurance models, noting the relatively disappointing prospects in the cyber insurance market: “Even a conservative forecast of 2002, which predicted a global market for cyber-insurance worth \$2.5 billion in 2005, turned out to be five times higher than the size of the market in 2008 (three years later). Overall, in relative terms, the market for cyber-insurance shrank as the Internet economy grew.” The authors note also that this weak cyber insurance market is often attributed to a lack of understanding: “The observable under-development of the market for cyber-insurance is often attributed to insurers’ lack of experience with a new kind of risk, combined with insufficient actuarial data hindering competitive pricing.” Thus, it is difficult for insurance providers to determine what insurance should cost, because of a lack of objective data informing the expected loss of a cyber event. This “lack of actuarial data,” as will be seen later, is not just missing in insurance models—it is indeed the missing piece in the cost/benefit model of cybersecurity investment.

B. MEASURING BENEFITS

1. The Gordon-Loeb Model

In 2002, University of Maryland Professors Lawrence Gordon and Martin Loeb developed a popular model for objectively evaluating investment benefit. The model takes an inductive, logical approach towards arriving at a final assessment of investment benefit. Essentially, the equation which Gordon and Loeb developed measures the economic benefit of information security as the difference between money spent on the investment itself, and the marginal decrease in expected loss due to the protection provided by the investment. Measuring the cost of the investment is relatively straightforward—it is simply the cost of whatever people, products, and technology are part of the cybersecurity investment. The impact of the investment, on the other hand, depends on the computed decrease in vulnerability. This is harder to measure inductively. Placing a baseline number on a system’s objective vulnerability is not easy, and no common, widely-used scale currently exists. However, the Gordon-Loeb model only asks

for the difference in vulnerability, before and after the security investment is made. That difference then informs the change in expected loss, assuming a static economic value assigned to the vulnerable asset. As put in their paper, “the key to analyzing information security decisions is not the vulnerability (or the expected loss without the investment), but the reduction in expected loss with the investment.” Thus, the Gordon-Loeb model elegantly simplifies the economic decision to one bit of information; the decrease in expected loss. This, however, is where the model fails; it plugs in arbitrary vulnerability functions without any real explanation as to why these relationships exist. In order to complete this model, true relationships between vulnerability and mitigation must be built and justified.

2. Empirical Analysis

Historical empirical analysis has been tried. Liu, Tanaka, and Matsuura present “a two-step empirical analysis of investing in security countermeasures based on a Japanese enterprise survey of security investments in the private sector (Liu, Tanaka, & Matsuura). The first step measured the vulnerability of Japanese enterprise by examining the change in reported computer virus incidents reported to the Japanese government between 2002 and 2003, and then examining which companies on both lists made significant cybersecurity investments. The study was made possible largely due to the mandatory reporting requirements of the Japanese Ministry of Economy, Trade, and Industry (METI), and survey data gathered from Japanese enterprises. The study found that significant investments in security protections (“countermeasures”) did indeed lead to a decrease in computer virus incidents. Persistent, or targeted threats may not have been affected. While only as good as the survey information provided, this study does provide an empirical test of the Gordon Loeb model and as such lays the groundwork for future study.

Another, more mathematical examination of the Gordon-Loeb model, performed by Jan Willemson (Willemson, 2006), frames the need for a further fleshing out of the relationship between investment and vulnerability: “In [their model], Gordon and Loeb considered two specific function families, but actually there is no reason to assume that any function in any of these families corresponds to any real vulnerability decrease

scenario. Thus, the main direction for further work is to look for functions reflecting changes in vulnerability for some real situations.” The “further work” to which Willemson is referring is precisely the work necessary to test cyber investment models, and thus inform decisions at the boardroom level.

Such a test of the Gordon-Loeb model would require an experimental framework, capable of isolating the variables inherent in a cybersecurity event, so that data points can be gathered, and possible correlations could be drawn. From this data, one may be able to build a continuous relationship between vulnerability and mitigation, specific to particular attack vectors. This relationship would fill out the Gordon-Loeb model, and given company-specific assessments of asset valuations, a more informed cost-benefit equation could be derived.

The investigation into the reasons behind cybersecurity investment thus begins with viewing the problem as the aggregate of independent, self-interested decisions by firms. From there, one can assess the various models which may apply—in this case, the Gordon-Loeb model. Understanding the firm’s decision process first requires assessment of the real cost of a breach. Additionally, one must examine non-investment options such as insurance. If investment is a choice, it must then be financially justified. Such a justification model exists in Gordon-Loeb, but it is missing a critical element: the relationship between vulnerability and cyber investment. However, this relationship is not currently possible to objectively ascertain with publicly available intrusion information. Moreover, it is highly likely that there are many such relationships specific to particular vulnerabilities and mitigations—as put by Willemson: “clearly, such functions are strongly application area specific.” Therefore, a plausible, and productive way forward in modeling private sector cyber investment decisions is to arm firms with objective, empirical data. If mathematical relationships between vulnerabilities and mitigations can be derived, that would further fill out the model. This is a critical missing piece in understanding the fundamental economic mechanics of the cyber ecosystem.

C. WHERE IS THE DATA?

Many people would agree that cyber security is a data-rich environment, but much of the real-world data resides in disparate corners of the private sector, academia, and can also possibly be classified by government agencies. Additionally, even if that data can be shared it is only useful in statistically significant amounts, and when it is normalized, formatted, and sensible.

This likely has much to do with the legal environment surrounding cyber incidents. While companies have become increasingly open about their intrusions in recent years, they do not share network diagrams, log statistics, or mitigation plans with the public, and for good reason. The details of a corporate breach involve proprietary and personal information, and as such remain protected. The process of sanitizing cyber attack data to exclude these sensitive details is very labor intensive, and organizations may not be able to justify dedicating resources to that effort, especially since it has little to do with addressing the problem at hand. Moreover, the tools and technologies deployed within an organization's network often belong to third party companies, and exposing the failure of those systems may similarly cause legal and contractual issues for companies. Most corporate press releases are limited to general dates and figures regarding the breach, as well as customer-designed guidance on protection from identity theft or fraudulent activities.

This state of affairs results in droves of anecdotes, third party stories, and rumors, but very little in the way of useful statistics. Not to mention, in many cases it is safe to assume that the victim company may still not fully understand the full nature of the breach, nor the preliminary measures taken by attackers. Therefore, even in the most verbose incident reporting, the data is still limited by the visibility and understanding of the reporting entity.

1. Existing Metrics

Several cyber metrics systems exist for recording real-world intrusion events—such as Verizon's VERIS framework (Verizon, 2014), the SANS Security Metrics (Payne, 2006), or the NIST standards for Incident Reporting (Cichonski, Millar, Grance, &

Scarfone, 2012), but very few standard metrics have yet been widely adopted to record and compare cyber actions taken during events. Therefore, it is nearly impossible for researchers at separate organizations to analyze the same data set from two different angles. With the exception of open-source projects like Project Honeypot (UnSpam Technologies, 2014), the HoneyNet Project (Project, 2014), and others, publicly available standardized cyber attack data is surprisingly scarce and rarely interoperable.

2. Exercises Hold Untapped Data

Cyber Exercises, especially those of a Red/Blue Team nature, have the potential to give researchers a “bird’s eye view” into a cyber attack. In almost all Red/ Blue team competitions, there is a White Team, or control team, which can change the situation on the fly to serve the objectives of the exercise. This team operates with omniscience and full awareness of the actions of both blue and red teams. If the data available to the white team could be recorded and indexed, it would provide the otherwise elusive visibility which does not exist in real-world intrusions, a perspective not subject to the biases or limits of either defender or attacker. The data would simply tell the story.

Additionally, the frequency and diversity of cyber exercises in recent years could potentially cover a wide variety of attack methods, mitigations, and technologies. If taken in aggregate over several years, just the range and frequency of cyber exercises could cover the most common attack methodologies. Since the attack and defense types are as diverse as the systems and configurations which comprise the cyber security ecosystem, it is also essential that simple, standard, technology-agnostic metrics be built which can apply to all cyber exercise situations. Similar to a batting average in baseball, these metrics must be universally usable and comparable across all exercises.

In order to yield useful information for real-world decisions, it is imperative that the exercise simulates a real-world situation as closely as possible, and that the statistics gathered are standardized and repeatable. The following section will examine cyber exercises and cyber attack methodologies, and propose criteria for the design of the proof-of-concept exercise.

THIS PAGE INTENTIONALLY LEFT BLANK

III. EXERCISE METHODOLOGY AND DESIGN CONSIDERATIONS

A. CYBER ATTACK METHODOLOGY

In the “Hacking Exposed” series of publications written by Stuart McClure, George Kurtz, and Joel Scambray, the “Anatomy of a Hack” chart appears on the inside of the back cover. Hacking Exposed 7th edition, for example, has the following steps, shown in Table 1.

1	Footprinting
2	Scanning
3	Enumeration
4	Gaining Access
5	Pilfering
6	Covering Tracks
7	Creating Back Doors
8	Denial of Service

Table 1. Hacking Exposed “Anatomy of a Hack” (from Scambray, McClure, & Kurtz, 2012)

While various iterations of this methodology have developed since Hacking Exposed’s First Edition, the “Anatomy of a Hack” provides a useful framework for discussing the phases of any cyber intrusion. More importantly for the purposes of this discussion, these modules can be used to group the myriad tactics and techniques within the cyber exercise universe. However, Hacking Exposed was written before much of the industry began focusing on the “Advanced Persistent Threat,” or “APT” style of attack, which often incorporates client-side exploitation and lateral movement (Symantec, 2014). Operations like these have been described in a collection of cyber research papers, from

the Mandiant APT1 report published in February 2013 (Mandiant, 2013), to the Elderwood Project paper published by Symantec in 2012 (O'Gorman & McDonald, 2012), back to the Shady RAT and Night Dragon reports published by McAfee in 2011 (McAfee, 2011) (Alperovitch, 2011). These types of intrusions center around movement and persistence within a compromised network, and eventual exfiltration of sensitive data. These operations can be mostly described using elements of the Hacking Exposed Anatomy.

To further supplement this categorization, however, the following life cycle included in the Mandiant APT1 report provides further modules and the idea of an internal loop in the methodology, shown in Table 2:

1	Initial Recon
2	Initial Compromise
3	Establish Foothold
4	Escalate Privileges
	Internal Recon
	Move Laterally
	Maintain Presence
5	Complete Mission

Table 2. Mandiant APT1 Methodology

Reportedly, many operations have targeted companies and organizations in the U.S. housing intellectual property central to their business model. Therefore, to provide data relevant to the problem noted in Chapter I, the design of a simulation exercise should follow the basic framework of these real-world APT intrusions. The next section briefly describes general categories of cyber exercises, and how to tailor an exercise more closely to an APT-style attack.

B. CYBER EXERCISE FRAMEWORKS

Cyber competitions have matured and diversified greatly in the last two decades. (Eller, 2004) Currently, most cyber exercises can be placed into four major categories, shown below:

Current Models of Cyber Exercises

- Disaster Response/ Coordination
- Capture the Flag- Race
- Capture the Flag - Battle Royale
- Vulnerability Discovery vs. Remediation
- Red vs. Blue

a. Disaster Response and Coordination

These cyber “exercises” most closely resemble physical disaster response exercises, and are often performed on a regional or national scale. The Cyber Storm series of cyber exercises, for example, follow this model (Department of Homeland Security, n.d.)

b. Capture the Flag–Race

These are Capture-The Flag competitions in which teams are given points based off how many “flags” they capture, on a neutral system. The event is timed and there is minimal inter-team activity. All teams are essentially attacking the same system, and the fastest team to capture all the flags wins. Points can also be deducted for hints or help, affecting the aggregate score (Symantec, 2014; The National Cyber League, 2013).

c. *Capture the Flag—Battle Royale*

In this scenario, teams try to capture flags which reside on other teams' networks, while defending their own. This incorporates both attack and defense, usually in a live-fire exercise. Figure 1 is an example of the “Defcon”-style exercise:




Figure 1. Defcon-Style Capture-the-Flag, (from Cowan, Arnold, Beattie, Wright, & Viega, 2003)

Another example of a Battle Royale-style exercise is the now-defunct Cipher CTF, which employed symmetric competition as well as functionality requirements. A description of the exercise, taken from the Cipher CTF website, is below:

The exercise consists of multiple teams, each hosting a server that has multiple services running, like e.g., a webserver, a mail server, or customized services. The services contain typical security vulnerabilities that allow to compromise the server to a certain extent. The goal is to maintain the services up, functional and uncompromised for the duration of the game. Additional scores can be gained by patching the vulnerabilities of the services and exploiting the knowledge of the found weaknesses at the other team's servers. (Pimendis, 2010)

It should be noted above that the Cipher CTF incorporates a functionality requirement, ensuring that network defenders do not lock down the network to the point of making it unusable. This is an important aspect of any real-world simulation, because it replicates—even if only to small extent—the pressures that network defenders deal with in securing their networks.

d. Vulnerability Discovery vs. Remediation

In some cyber defense competitions, points are awarded to teams that can identify and remediate vulnerabilities first, and provide a proof that the vulnerability is fixed. While this is an important aspect of blue-team system and network hardening prior to an incident, it reveals little about incident response capability (Air Force Association, 2013) (DARPA, 2014).

e. Red vs. Blue

In this model, one red team is tasked with attacking one blue team network. The goal is often the access to or exfiltration of sensitive data. Points are often awarded to the Red team for certain “flags,” or tasks they have to accomplish. While some “hacking back” may be seen from the blue team, the blue team’s main objective is simply to hold off the red team for as long as possible until the exercise ends.

Therefore, the exercise type that most closely simulates APT activity against the U.S. private sector and civilian government is the Red/Blue design, for several reasons.

Active Threat—Most Red/ Blue team exercises involve human beings on both sides—hence a continual active threat, and active, aware defense. This most closely resembles the conditions of APT intrusions, where malicious actions are thought to be mostly done manually.

Functionality Requirements—In a Red/ Blue team exercise, the Blue Team is defending a simulated “real network.” This often involves networks with an expected functionality, not simply a test range. Thus, various services and protocols must be open, for business to occur. This also closely resembles many private and public sector networks, where functionality often trumps security.

Exfiltration Goal—Finally, Red vs. Blue Team exercises often involve a Confidentiality or Integrity compromise—the type of compromise about which many companies are concerned when it comes to APT actors. That goal determines much of the pace and mitigations steps of a blue team, versus, say destructive malware or DDoS.

C. CONDITIONS OF THE PROPOSED EXERCISE FRAMEWORK

The “Storyline” of the proposed exercise framework follows a blend of the Hacking Exposed and Mandiant APT1 methodologies, with discrete steps toward achievement of exfiltration of sensitive data. Additionally, the exercise should be constructed to allow for the Red Team to succeed initial compromise, with minimal, but necessary, assistance by the White cell. This helps the exercise to test cyber defense capabilities, and not solely pre-configuration or defensive hardening abilities on the part of the blue team.

a. Baseline User Activity/ Traffic

Real-World networks are dominated by “civilians”—meaning it is not just red and blue team members. Much of cyber defense is finding malicious activity within large amounts of “legitimate” traffic. Having White Team members pretend to be legitimate “users” on the network is one way to accomplish this. There are several open-source network traffic simulators available, and these could be used to create the signal/noise ratio in the exercise environment.

b. Intelligence/ Threat Information

Prior to any attack, real-world cyber defense teams may likely have some understanding of the threats they face—actors and attack vectors. Beginning any exercise with some basic threat intelligence would help simulate real-world scenarios and give the blue team a way to prioritize resource allocation.

c. Blue Team Resources/ Functionality Requirements

During an attack, some have argued that the blue team should just unplug the network. In real-world scenarios, this is rarely an option. Blue teams have to keep the network to an acceptable level of functionality, or at least pay a penalty for degraded performance for legitimate users. This was incorporated, for example, into the above-mentioned Cipher CTF exercise. Additionally, the blue team cannot have unlimited monitoring, storage, or analysis resources. These constraints will then help emulate the

real-world decisions that cyber defenders face—decisions which are the very foundation of effective defense and incident response.

d. Preexisting Configurations

Few cyber defense teams get to build their systems from the ground up. If possible, blue teams should defend systems from an existing state while keeping core functionalities running. This best simulates real-world scenarios where security analysts and leaders inherit the existing vulnerabilities on their network, and work to fix those as best they can. Even more so for cyber-incident response teams, who may know very little about the customer network before asked to come on site.

e. Real-World Duration

Cyber defenders are already busy enough. It is hard to find time to do training, and thus many exercises are constrained to within a week or two. Given this real-world time constraint, it can be difficult to simulate an APT campaign, or even a criminal cyber attack, which may in reality happen over months—or years in some APT cases.

f. Scope

Depending on the exercise architect, the contested network must be scoped. While large-scale intrusions may take place on networks with thousands of machines, exercises typically operate at several orders of magnitude lower. To replicate the internal network of even a medium-sized corporation would require hundreds of machines, licenses, and infrastructure devices, which most classroom environments may not have. Instead, exercises should seek to use representations; a handful of desktop/laptop machines, for example, can effectively represent a 1000+ machine user environment.

g. Provide Avenue for Initial Compromise

With an advanced and persistent threat, eventual initial compromise of some asset is all but guaranteed. Sooner or later, a vulnerable system (known or unknown) will be exploited. A user will click on a link or open an attachment. A USB drive will be plugged into a system. A vulnerable web server will be promoted into production, or a zero-day

exploit will bypass even the best security team. Focusing on prevention is futile—what is important is early detection, resilience, and eradication before sensitive data can be stolen.

h. Provide Avenues to Move Laterally/ Establish Persistence

Once inside the network, APT actors will often utilize preexisting infrastructure to move laterally –using legitimate accounts, file shares, etc. These methods are precisely the same tools and technologies that legitimate users utilize, and as such they must remain enabled for the exercise. This speaks to the blue team’s baseline functionality requirements; they cannot simply shut down the basic functions of the network, but rather must find the anomalous behavior in logs and artifacts.

i. Timing

To allow for full activities on both sides, but also a full ability to log and record results, an exercise could contain three phases, each entailing several steps. An example phase breakout is below in Table 3, using the various modules from the Hacking Exposed Anatomy and the APT1 methodology:

Phase 1	Scanning
	Enumeration
Phase 2	Establish Foothold
	Escalate Privileges
	Maintain Presence
	Move Laterally
Phase 3	Complete Mission
	Cover Tracks

Table 3. Phases of Offense

If conditions such as those noted above can be created in a classroom environment, then the actions and reactions of an exercise in that environment can provide a reasonable simulation for actual cyber incidents. The following chapter will describe the topological, technical, and temporal aspects of a, localized, easy to replicate “One Man” exercise, as well as a classroom-setting Red vs. Blue Exercise, both of which

should resemble real-world intrusions. These will not only provide models for educating cyber defenders, but also for gathering statistics for use in cyber modeling efforts. These statistics will need to be standardized, repeatable, and granular enough to allow for multiple combinations and derivative statistics in future studies.

D. STATISTICAL MEASURES

1. Indicators and Infrastructure

The statistics are intended to measure how effective the Blue Team's detection efforts were against the Red Team's infrastructure. These can be broken down by attack phase and totaled at the end. Table 4 breaks these team-based measures and statistics down.

Measures	Data Type
Infrastructure Nodes Used by Red Team	Count
Unique Network Infrastructure Nodes Enumerated by Blue Team	Count
False Positive Network Indicators	Count
Artifacts Placed on Blue Network by Red Team	Count
Artifacts Detected by Blue Team	Count
False Positive Artifacts Found by Blue Team	Count
Accounts Available to Red Team	Count
Accounts Used by Red Team	Count
Compromised Accounts Discovered by Blue Team	Count

Statistics	Data Type
Red Team Infrastructure /Blue Team Indicators	Ratio
Network Nodes Signal/Noise	Ratio
Artifacts Used Versus Detected	Ratio
Artifact Signal/Noise	Ratio
Accounts Used Versus Discovered	Ratio
Account Signal/Noise	Ratio
"Legitimate" Tools Used by Red Team	Number

Table 4 Team-Based Measures and Statistics

2. Tool-Based Statistics

These statistics are meant for network defenders to evaluate which tools helped the most during a given exercise. Over the course of many exercises, these aggregate statistics will show which tools are usually among the most effective for network defenders. These stats can also be broken down by phase and tabulated at the end.

Measures	Data Type
Number of True Positives	Count
Unique True Indicators	Count
Number of False Positives	Count
Number of Instances of Tool	Count
Red Team Blocks	Count

Statistics	Data Type
Signal/Noise Ratio	Ratio
True Positives Per Instance	Ratio
Red Team “Blocks”	Number
% of All True Indicators	Ratio
% of All True Pre-Exfiltration Indicators	Ratio

Table 5. Tool-Based Measures and Statistics

E. INCORPORATING STATISTICS INTO CYBER INVESTMENT MODEL

Hypothetically, for every APT exfiltration attack, there is a probability, at any given time, that it will either succeed or not. The job of the security administrator is to minimize the probability of a successful exfiltration. To do so, the security team uses detections and mitigations.

1. Evaluating Alerting Technologies – Theory

Whether at the network or host level, cyber defense technologies exist to detect, alert, and block badness from occurring. Ideally, a cyber defense technology would detect all malicious indicators within its purview, and only the malicious ones (zero false

positives). Additionally, the technology would not miss any malicious activities (zero false negatives). In order to circumvent this, attackers seek to bypass, obfuscate, or change their infrastructure.

Rather than trying to build an objective measure of benefit from the ground up, this approach allows an aggregate of attackers to empirically determine which technologies and procedures have the greatest impact to their operations. The benefit of a security asset is then by definition its negative effect on offensive operations.

For an attacker to make steps forward in their operation, they must use infrastructure. Infrastructure has a chance of setting off alerts, which can provide indicators to the blue team. The blue team then can remove some percentage of the necessary infrastructure from the red team. In order to fully prevent the red team from moving forward (a “block”) the blue team must detect and mitigate all infrastructure in that step, assuming perfect detection (no false positives) and full infrastructure enumeration (all infrastructure identified).

Theoretically, an attacker has a 0% chance of achieving his or her goal if all their offensive infrastructure is discovered. This includes external IP address space they may use as a jump point, and malware and exploits at their disposal, but also the “internal infrastructure” gained after initial compromise. This includes stolen account credentials and internal points of presence, for example. As a cyber defender, it is crucial to find and eliminate this infrastructure, or at least take it away from the attacker. Thus, the share of prevention belonging to any one technology can be reasonably measured by the *percentage* of actor infrastructure that a given technology is responsible for finding. As will be shown later, this approach can also be tailored to penalize for false positives and false negatives. Measuring the effectiveness of a technology is, in mathematical terms, to measure the intersection between Set A (actor threat infrastructure) and Set B (all technology-discovered indicators). When applied to a controlled exercise, the expected benefit of a technology can be viewed as the sum of all infrastructure enumerated during the exercise. This is the net impact that this technology could have potentially had on the adversary, assuming a perfect security team. The next section breaks down the terms and calculation of this metric.

2. Properties of Indicators

For each indicator “I” provided by a given technology, there are the following properties:

1. Appearances of “I” in total number of records generated by technology, and;
2. Percentage of similar infrastructure available to adversary.

For an example of the first, take an IP address detected by an IDS during a scan. The IP address appears in 5% of the alerts. While the technology obviously believes something is awry, most alerts require validation and examination for false positives. In order to evaluate the technology independently of human beings or skillset, this metric assumes a random choice of any alert as an indicator. In this case, there is a 5% chance that this indicator would be randomly chosen out of all alerts.

With regards to the second, it is also important to properly weight the relative importance of a given indicator. For example, if the network administrator blocks an IP address, but the actor has 10 similar IPs available from which they can scan, the administrator has only eliminated 10% of the problem. Thus, the expected benefit of this indicator is the chance of discovering it in logs times the marginal impact it has on the active adversary infrastructure. As another example, if a malicious file is found on a host machine, but there are 10 such infections, the administrator has only removed 10% of the problem.

3. Implementation Into a Simplified Gordon-Loeb Model

Remember that the Gordon-Loeb model makes the reasonable assertion that Expected Benefit of Information Security (EBIS) Investment is equal to the expected change in loss due to a breach. Assuming an APT-style data exfiltration scenario, the cost of a data breach can be estimated by the intellectual property valuation of the data that stands to be stolen. That is:

$$\text{EBIS} = P(s) * \text{Worth of Asset}$$

Equation 1. Simplified Gordon-Loeb

This P(S) can be substituted by a metric called “Detection Impact.” For a given technology, Detection Impact is measured as the signal to noise ratio of the technology combined with the impact of the indicators the technology discovers.

What this looks like, mathematically, is:

4. Detection Impact

For all Indicators I_1, I_2, I_3 , etc generated by a given Technology T, the technology has an aggregate impact on the adversary’s ability to complete their mission. If there are three technologies TA, TB, and TC, each responsible for discrete, mutually exclusive indicators:

$T_A: I_{A1}, I_{A2}, I_{A3}$

$T_B: I_{B1}, I_{B2}, I_{B3}, I_{B4}$

$T_C: I_{C1}, I_{C2}$

Then the Detection Impact of Technology A, represented by $D(T_A)$ is equal to the below:

$$D(T_A) = \sum_n (\text{Signal to Noise Ratio of Indicator } I_{An}) \times (\text{Indicator Weight Factor})$$

Equation 2. Detection Impact

Signal to Noise Ratio for Each Indicator” is calculated as:

$$\frac{\text{True Positive Alerts Generated Conataining Indicator } I}{\text{All Alerts Generated by Technology } T}$$

Equation 3. Indicator Signal To Noise Ratio

The Weighting Factor represents the share of like actor infrastructure that an indicator represents. For example, if the attackers have 3 exploits available, and the technology finds 1, that weighting factor is 1/3.

$$\frac{1}{(\text{Number of Like Infrastructure at Attacker's Disposal})}$$

Equation 4 Indicator Weight Factor

$$\sum \left[\frac{(\text{Alerts Generated containing Indicator } I)}{(\text{Total Number of Alerts Generated By Tool})} \right] \times \left[\frac{1}{(\text{Number of Like Infrastructure at Attacker's Disposal})} \right]$$

Equation 5 Breakout of Detection Impact

As an example of how the summation works, suppose a technology detects three indicators. The calculation of the Detection Impact is in Table 6.

	Indicator 1	Indicator 2	Indicator 3	Total
True Positive	20	15	17	52
Total Alerts	100	100	100	100
Signal/Noise Ratio	0.20	0.15	0.17	0.52
Indicator	1	1	1	3
Total Infrastructure	25	25	25	25
Weighting Factor	0.0039	0.0039	0.0039	0.0039
Technology Detection Share	0.1%	0.1%	0.1%	0.2%

Table 6. Calculation of Detection Impact

Table 7 is an example Weighting Factor Calculation, given an actor with a /24 subnet, 20 owned internal machines, and 2 accounts enumerated.

Weight of Indicator			
	IPs	Machines	Accounts
Available	256	20	2
Weighting Factor	0.004	0.025	0.500

Table 7. Calculation of Weighting Factor

5. Summary

The Detection Impact metric rewards high true positives and penalizes false negatives, allowing defenders to compare technologies across various levels of the security stack. For example, an IDS detecting 50% true positives, but which enumerates 20% of available actor IP space, has a $50\% \times 20\% = 10\%$ impact on the actor's ability to complete their mission. Conversely, a host-based AV with 90% true positive rate, but which misses 1 of 10 malicious artifacts has a $90\% \times 10\% = 9\%$ impact on the actor's operation. This metric can also inform cyber defenders where they can optimize their impact on the attackers—for example, using IP lists to block APT activity, or playing “whack-a-mole” with an IP block list, will not have much impact. However, hardening hosts and updating defensive signatures could potentially have a significant effect on an actor using commodity malware, by preventing or alerting on 100% of infected systems. This would present a high yield mitigation step in the cyber kill chain. Thus, in addition to evaluating the tool's detection abilities, this metric also rewards tools that can produce high-value indicators.

THIS PAGE INTENTIONALLY LEFT BLANK

IV. METHODOLOGY AND DATA GATHERING PROOF OF CONCEPT

A. FRAMEWORK FOR ONE OR TWO-PLAYER EXERCISE

Overview:

In order to demonstrate how these metrics are measured, I created a small battle environment to conduct an exercise. The steps in this exercise fall into common methodology steps, and the network under attack retained its functionality throughout the exercise, hence satisfying the requirements listed in the previous chapter. This exercise modeled the exfiltration of data from an internal Windows environment. It incorporates scanning, exploitation—both outside in and client side—password cracking, lateral movement, and exfiltration.

For the purposes of this exercise, each machine involved represents a different “segment” of an enterprise network:

Host/ User Machine: Internal network

Web Server: DMZ

Windows 2003 Server: Enterprise Infrastructure

Kali Linux: Evil Internet

Security Onion: Security Infrastructure (Security Onion, 2014)¹

Assumptions:

Footprinting: Because this was a single-user exercise, executed by myself, I bypassed the “Footprinting” stage. Since I built the entire exercise, it was impossible to not know the configurations or details of the blue network. A proper Footprinting step is outlined in the classroom exercise description.

Flat Network:

While it is possible to include a virtual router and/or firewall into the single user exercise, this exercise assumes a “flat” network. Because each machine in this exercise is representative of an enterprise segment, there was no functional need for a router, as shown in Figure 2.




Figure 2. Single or Dual User Exercise Topology

Scenario Design

Offensive Steps:

1. Scan Network
2. Exploit Web Server
3. Client Side Exploit of Windows XP machine
4. Hashdump/ Password crack to get Domain credentials
5. Move laterally to internal Windows server
6. Exfiltrate data
7. Cover Tracks

Defensive Functionality Requirements:

1. Keep all machines up
2. Freeze IDS and AV signature updates
3. Cannot block subnets

Hardware/Software Used

I used three physical computers, mostly due to storage and performance concerns. I connected them with a hub and CAT5 Ethernet cables, and kept them air-gapped from any other network connections.

Hardware: Dell Laptop running Host OS of Kali Linux (Offensive Security Ltd., 2014)²

Dell Laptop running Virtualized windows environment (Internal network)

- Windows Server (DNS/ Active Directory/ Email)
- Windows Host

Dell laptop running virtualized Linux environment (DMZ)

- Metasploitable (Rapid7, 2012) (web server)
- Security Onion (IDS)
- Ethernet hub with 4 network interfaces
- 3 CAT5 Ethernet cables

Defensive Tools:

- Intrusion Detection System - Security Onion
- Network Traffic: Wireshark
- Host Anti Virus—Avast Free AV
- Host Event Logs—Windows Events, Linux var/log

Offensive Tools:

- Main Interface: Armitage/ Metasploit
- Exploit Package Development: Social Engineering Toolkit
- Scanner: Nmap
- Password Cracker: John the Ripper (Johnny)

B. PROCESS/ RESULTS

Using the Kali Linux machine, I began a Metasploit database and started Armitage. This allowed for a graphic depiction of the attack, as well as a good place to consolidate scans, sessions, and data all in one place, as shown in Figure 3.




Figure 3. Mapping the Network

Phase 1: Scanning:

First step was to scan the /24 network housing the webserver and other machines in the network. For this I used Nmap within Armitage.

Enumeration: Upon scanning the Metasploitable web server, I (as expected) found several vulnerable programs. Armitage allows for a “Find Attacks” option, which examines the scan information for possible exploits.

***Note:** The exploitation of a Metasploitable server is neither difficult nor unique in nature. The use of this well-known vulnerable VM was to facilitate the exercise, not to demonstrate offensive expertise. Unfortunately, real-world anecdotes often reveal systems similarly vulnerable, so it is not out of the question that such a server might be running in the wild.*

Phase 2: Exploitation:

Using Armitage, I then executed a “Hail Mary” against the server, to see what exploits would work and which sessions I could begin. Sure enough, 4 sessions opened,




Figure 4. Sessions Started on Web Server

with varying privileges, as seen in Figure 4.

Lateral Movement: Once a session was established with the Metasploitable server, I sought to make my next move in the methodology—move laterally to the windows domain. To do this, I created an embedded .pdf executable using the Social Engineering Toolkit (SET): Once the executable was constructed, I used a session from Metasploit to upload the executable to the `var/www/` directory of the Metasploitable web server, as shown in Figure 5.

Client-Side Exploitation (Maintain Access): The next step assumes a social engineering success to have someone in the Windows domain access the pdf on the Metasploitable web server. I felt comfortable making this assumption because this is not only a common social engineering tactic, but in this scenario the web server and Windows hosts were part of the same organization.




Figure 5. Uploading the Exploit PDF

On the Kali Linux machine, I began a listener on port 443. I coded this port as a callback on the pdf executable because it is commonly allowed out on most internal networks. Once the pdf was downloaded by a victim on the web server, it immediately executes and calls back to the attacking computer. Now that I had access to the windows domain, I performed an “arp -a” command on the XP host to see what other computers it was talking to. The arp sweep revealed the Windows server.

Moving Laterally: Additionally, while on the XP machine, I performed a dump of all the stored credentials, using the “hashdump” meterpreter command. Once I had the hashes, I set about cracking them using John the Ripper with the default wordlist that came with the Kali Linux build. Once the Administrator password had been cracked, I sought to map to a share within the Windows domain to simulate lateral movement. I used the net view/ net use commands, authenticating to the Windows server as Administrator. This is shown in Figure 6.

Phase 3: Exfiltration:

Once I had successfully mapped the drive, I could explore the file system through meterpreter, and download the secret file, as shown in Figure 7.

Cover Tracks: Finally, to cover my tracks, I used the “clear event log” function within the meterpreter shell, on the XP host. Additionally, I wrote “” to the var/log/syslog file on the Metasploitable server. I was unable to wipe logs on the Windows server, however, since I did not have shell access but rather simply mapped the drive.




Figure 6. Mapping Network Drive




Figure 7. View/ Exfiltrate Sensitive Information

C. MONITORING/ GATHERING STATISTICS

On the blue team side, there were several places to gather both live and after-the-fact statistics on the attack. I wanted to be sure that these tools would indeed gather alerts, and also examine the data to indeed see if the statistics were viable and sensible.

Network Technology:

The security onion VM served as a useful, if somewhat finicky, IDS alert platform. Security Onion provides a wealth of different IDS and PCAP technologies; I chose to use SQUIL in this instance, as shown in Figure 8.




Figure 8. IDS Alerts

Network IDS Technology Evaluation:

The following calculations were performed to calculate the *Detection Impact* of the IDS technology. There were 97 total alerts from scanning, revealing the attacker IP address. Additionally, a total of 6 alerts fired on exploitation, covering 2 of the 4 exploits launched.

	IP Address 1	Exploit 1	Exploit 2	Total
True Positive	97	4	2	101
Total Alerts	174	174	174	174
Detection Rate	0.56	0.02	0.01	0.58
Indicator	1	1	1	
Total Infrastructure	256	4	4	
Weighting Factor	0.0039	0.2500	0.2500	0.1680
Detection Impact	0.2%	0.6%	0.3%	9.7%

Table 8. Calculation of IDS Detection Impact

Host- Based Technology Evaluation:

I installed Avast! Free Anti-Virus on the XP host. While it alerted on the initial pdf exploit and another malicious .pdf file I uploaded, (as seen in Figure 9), it did not alert on a persistence agent uploaded after a session had been established. This speaks to the low false positive rate of Anti-Virus solution—but also the false negative rate. In fact, Symantec recently reported that they estimate their AV solution only catches 45% of current cyber attacks (Yadron, 2014). In comparison to the IDS, the Detection Impact statistic accounts for the false negative rate, even though the Signal/Noise ratio was perfect.

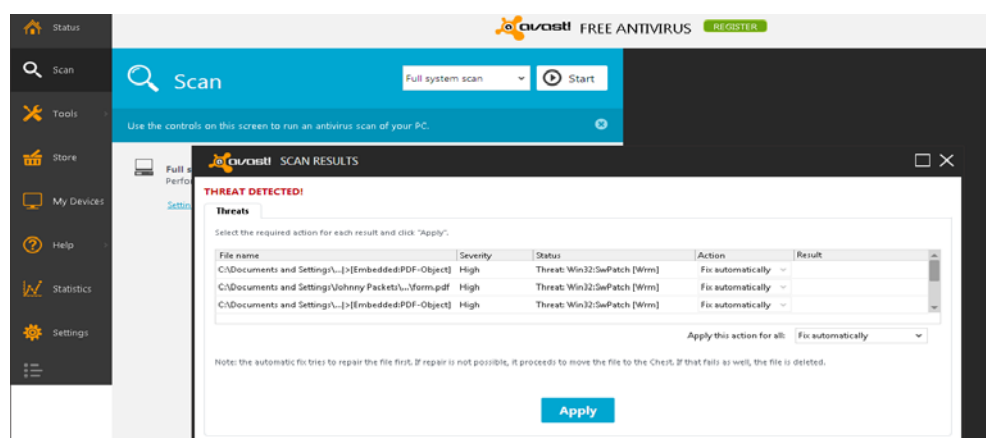


Figure 9. Avast! AV Alerts

	Artifact 1	Artifact 2	Artifact 3	Total
True Positive	4	3	0	4
Total Alerts	7	7	7	7
Detection Rate	0.57	0.43	0.00	0.57
Unique Indicator	1	1	1	
Total Infrastructure Type	3	3	3	
Weighting Factor	0.3333	0.3333	0.3333	0.3333
Detection Impact	19.0%	14.3%	0.0%	33.3%

Table 9. Calculation of Avast Detection Impact

To elaborate on this metric, the reason why the AV solution had a higher Detection Impact is primarily because had AV alerted on all three artifacts, and assuming rapid or immediate mitigation had been taken against them, I as the adversary would have been effectively blocked from my mission.

Summary:

The Detection Impact metric was intended to provide a platform-agnostic measure of a technology's ability to provide a blue team with a full, clear picture of malicious actor's presence on the network. While the metric may not account for technological externalities or extreme cases, it can be used to compare detection technologies at the

network, domain, and host levels. If the aggregate score of a security stack approaches 1, then the Blue Team is in a position to effectively block the actors from success in their mission. In the examples of the Single-User Exercise, the Detection Impact metric did indeed land between 0 and 1 for both IDS and AV, and provided comparable values of 9.7% and 33.3%, respectively.

D. EXTENSIONS TO CLASSROOM EXERCISE

Building off the basics of the small exercise, a classroom exercise can be built with similar principles, with more machines, more technologies, and thus more opportunities for evaluation. Instead of single machines representing the Attacker, DMZ, and Internal Network, more machines can be placed into segments, better simulating a real enterprise network. An example of what a classroom exercise might look like is in Figure 10.

Overview:




Figure 10. Classroom Exercise Topology

Essentially, the exercise should preserve the concepts and methodology of initial compromise, privilege escalation, lateral movement, and exfiltration. This methodology resembles real-world activity, and will also help make sense of the statistical data gathered. If the exercise were extended to the classroom environment and with a larger network, the methodology could look like the below:

1. Scan Network
 - a. Can scan the whole network, or just DMZ. Additionally, Red Team can switch/spoof IPs
2. Exploit Web Server

- a. There should be several servers in the DMZ. Web server could form an initial point of entry, but DNS and/or Email could also be leveraged.
3. Client Side Exploit of Windows XP machine
 - a. This could take several forms. In my example, I pushed a malicious file to the webserver—but a spear-phishing attack or other abuse of trusted relationships between the DMZ and internal network could be used.
4. Hashdump/ Password crack to get Domain credentials
 - a. This could entail compromise of a domain controller, rather than simple password cracking; however, cracking hashes or pass-the-hash remain relatively simple ways to achieve the same goal.
5. Move laterally to internal Windows server
 - a. Since there are several servers, this can be expanded to include a file server or database.
6. Exfiltrate data
 - a. Rather than a direct download through a shell, a classroom exercise could incorporate an internal blue network FTP server as a pivot point.
7. Cover Tracks—This could potentially involve data deletion, or centralized logging corruption—depending on the scope of the exercise.

Timing:

The three phased approach in the single/dual user exercise worked well, allowing me to perform Red Team functions, stop, and then examine the Blue Team detection tools for alerts and information. Therefore, the proposed timing of the classroom exercise is to break the exercise into three phases as well. However, these phases are further split

into “innings,” alternating activity between teams in order to objectively measure the effectiveness of the team, tools, and tactics.

	Phase 1	Phase 2	Phase 3
Offense	<i>Scan</i>	<i>Gain Privileges</i>	<i>Complete Mission</i>
	<i>Enumerate/ Exploit</i>	<i>Move Laterally</i>	<i>Cover Tracks</i>
	Submit Vulnerable Hosts and Services, along with Infrastructure used, to White Cell	Submit Internal Network Infrastructure to White Cell	Submit Exfiltrated Info to White Cell
Defense	<i>Examine IDS Alerts and FW Logs</i>	<i>Examine AV and other Network/Host Technologies</i>	<i>Examine Forensic Images</i>
	Submit IP and signature Indicators to White Cell	Submit Artifacts and Accounts to White Cell	Submit Exfiltrated Info and Evidence to White Cell

Table 10. Classroom Exercise Phases

Hacking Back:

For the purposes of this exercise, “hacking back,” that is, offensive activities undertaken by the blue team, are not allowed. The reason in this case is that in a “hack back” scenario, the offensive blue team’s objective is not data exfiltration but rather destruction and disruption. While in a wartime scenario, this may apply, in economic espionage cases it represents a fundamentally separate paradigm.

Additional Modules to Consider:

Security Team Under Surveillance: In Phase 2, the Red team gains visibility on the security devices and actions of the Blue team.

Destructive Payload: In Phase 3, the red team destroys or deletes data as a final step, rather than exfiltration. It is important for assessment purposes that records be copied by the white team before this action, in order to ensure historical data is not deleted as part of the exercise itself.

DDoS as a distraction: During any of the phases, DDoS could be used by the Red Team to distract or disrupt Blue Team operations in the midst of the exercise.

THIS PAGE INTENTIONALLY LEFT BLANK

V. CONCLUSIONS AND FUTURE WORK

A. LESSONS LEARNED

Executing the single-user/ dual user exercise, and recording statistics while doing so, illuminated many lessons about how statistics might be gathered in a classroom setting. Additionally, the various different iterations of systems that needed to be created, destroyed, and created again made it clear that these statistics should be general, not specific to any given technology. They must also be specific enough to highlight attacker and defensive weaknesses at every stage of the exercise.

Given the technical effort involved in full statistical capture for the single-user exercise, the classroom exercise could become prohibitively complicated if not properly scoped and planned. In order to gather meaningful statistics, there must be a solid understanding of security device configuration, visibility, and output in order to objectively measure its effectiveness. Additionally, while maintaining an authoritative timeline and a compendium of statistics is a hard task, it must be a priority for the white cell.

Cyber defense technologies often involve complicated and obscure commands, interfaces, and terms. While it may be easy to blame the opaqueness of a tool's output on the user, there is something to be said for ease of usability and interpretation. If a tool requires months or years of training, then the investment for that tool is far higher than the sticker price. Therefore, while missed alerts or indicators may technically be due to a blue team's lack of skill, tools which require extensive training will be scored lower than easy-to-use tools.

B. FUTURE WORK

While measuring a technology's ability to detect actor tools and infrastructure is a fair measure, that knowledge is no good if a blue team cannot remove malicious access in time. Statistics must be tested to measure the time it takes for a blue team to translate alert data into indicators, and indicators into mitigations.

Time-Based Statistics:

Theory: If a mitigation action is taken to remove actor access after the actor has already achieved their mission, there is little point. The important thing from a cyber defense perspective is to operate inside the operation loop of the adversary. Thus, the grand measure of success or failure is whether or not alerts yielded indicators, and indicators precipitated mitigations, fast enough to affect the adversary's goals.

Implementation: Each phase of the exercise is evaluated on time of first alert, to time of first indicator submitted, to time of first mitigation put in place. This measures people and processes ability to analyze alert information, gain confidence in indicators, and then implement mitigations. Thus, the time from first alert to mitigation, as compared to first contact to mission completion, is the measure of the effectiveness of a process or personnel during an exercise.

The below statistics could potentially be used to measure the team's timely or untimely performance—measuring reaction to information, ability to adapt, quality of detection, and other qualities which apply to the human components of red and blue teams. These metrics require testing and validation, as part of future projects.

Measures	Data Type
Time of First Contact—White Cell	Time
Time of First Alert (IDS or AV)	Time
Time of First Compromise	Time
Time of First Exfiltration	Time
Time of Final Packet	Time

Statistics	Data Type
Time until Network Compromised (First Entry)	Duration
Time between First Contact and First Alert	Duration
Time between First Contact and Compromise	Duration
Time between Compromise and Exfiltration	Duration
Time until Data Exfiltrated	Duration
Time Between First Alert and Exfiltration	Duration

Table 11. Time-Based Statistics

Mitigation- Based Statistics:

It is imperative that the blue team keep track of what mitigations are put in place, and when. This is critical to recognizing which mitigations had the most effect on the actors' ability to exfiltrate data.

Metric	Data Type
Network Mitigation:	(Time, Indicator)
Artifact Mitigation:	(Time, Indicator)
Account Mitigation:	(Time, Indicator)

Table 12. Mitigation Statistics

For each mitigation (Indicator Submitter, Layer of Mitigation, Time of Submission), the timeliness can be viewed using the following approach

Mitigation Timeliness: For Each Mitigations M which is informed by Alerts A₁, A₂, A₃, etc containing Indicators I₁, I₂, I₃, etc:

Average [(Time from First Alert A1 to Mitigation M1—(Time from First Contact to Exfiltration))] / (Time from First Contact to Exfiltration)

Equation 6. Mitigation Timeliness

This concept is portrayed graphically in Figure 11.

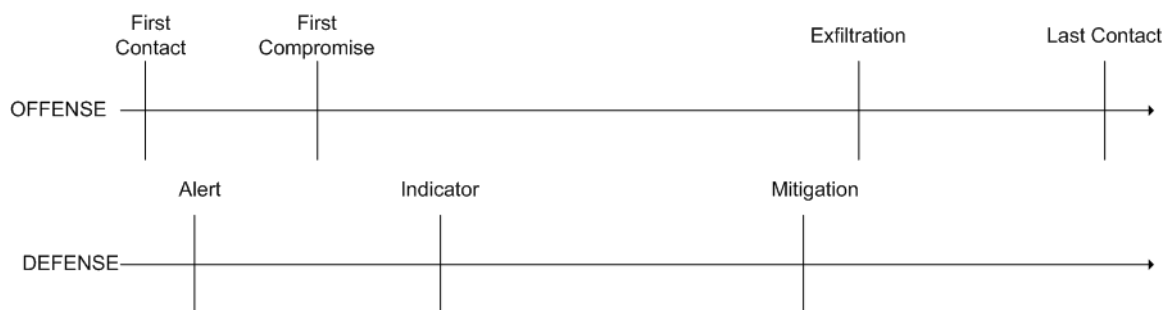


Figure 11. Mitigation Timeline

Rather than simply evaluating technology, these metrics can be used to evaluate processes and technology—namely, room for improvement. That can justify a need to devote resources to procedure development and practice, or training for Incident Response Personnel to more quickly turn alerts to indicators and indicators to mitigations. While not as discrete as the metrics for technology implementation, they can provide a good measure of which processes are most important to focus on, or which areas of a team or procedure need improvement.

C. CONCLUSIONS

In order for the concept of exercises as sources of statistics to help cyber modeling efforts, cyber exercises need to adopt more robust statistical gathering mechanisms, and share them openly with the community. This paper has demonstrated the need for and means by which these statistics can be generated and integrated into existing models. These are first attempts; it is my hope that better and more elegant manners of translation are built and executed.

Secondly, it is imperative that there be a clearinghouse of these statistics, and that they be made public for everyone's consumption. These statistics will gain accuracy as the sample size grows. It is possible, for example, to publish the single user exercise and let people submit their own statistics with verification from a third party. If the data is crowd-sourced in that manner, it could gain traction in the community and help to build a

much larger data set of attack and defense stats. Additionally, if classroom Red/Blue exercises track statistics properly, that data can be incorporated into the larger community. If properly indexed and categorized, this data could finally solve the lack of operational data that cyber security research suffers from.

Finally, cyber security researchers must use this data to build and test cost and investment models. If these models can be developed, tested, and used to justify real world decisions.

THIS PAGE INTENTIONALLY LEFT BLANK

LIST OF REFERENCES

- Air Force Association. (2013). *CyberPatriot rules book*. (Northrop Grumman) Retrieved June 6, 2014, from <https://www.uscyberpatriot.org/competition/rules-book>
- Alperovitch, D. (2011, August 2). *Revealed: Operation shady RAT*. Retrieved June 6, 2014, from <http://www.mcafee.com/us/resources/white-papers/wp-operation-shady-rat.pdf>
- Anderson, R. (2001). Why information security is hard—an economic perspective. *17th Annual Computer Security Applications Conference*. New Orleans, LA: University of Cambridge Computer Laboratory.
- Bohme, R., & Schwartz, G. (2010). Modeling cyber-insurance: towards a unifying framework. *The Ninth Workshop on the Economics of Information Security*. Harvard University.
- Cichonski, P., Millar, T., Grance, T., & Scarfone, K. (2012). *Computer security incident handling guide*. Gaithersburg, MD: NIST.
- Cowan, C., Arnold, S., Beattie, S., Wright, C., & Viega, J. (2003). Defcon capture the flag: defending vulnerable code from intense attack. *DARPA DISCEX III Conference*. Washington, DC.
- DARPA, D. A. (2014, May 16). *Cyber grand challenge*. Retrieved June 6, 2014, from https://cgc.darpa.mil/CGC_Rules_16_May_14_Version_2.pdf
- Department of Homeland Security. (n.d.). *Cyber Storm: securing cyber space*. (DHS) Retrieved June 6, 2014, from <http://www.dhs.gov/cyber-storm-securing-cyber-space>
- Eller, R. (2004, October 15). *Black Hat Japan 2004 - capture the flag games/ measuring skill with hacking contests*. Retrieved June 6, 2014, from <http://www.blackhat.com/presentations/bh-asia-04/bh-jp-04-pdfs/bh-jp-04-eller/bh-jp-04-eller.pdf>
- Ferrara, E. (2002). *Determine the business value of an effective security program - information security economics 101*. Forrester Research.
- Global Cyberlympics. (2014). *Global Cyberlympics*. (Global Cyberlympics) Retrieved June 6, 2014, from http://cyberlympics.org/?page_id=722
- Kataria, R. B. (2006). “Models and measures for correlation in cyber-insurance.” *Fifth Workshop on the Economics of Information Security*. Cambridge, UK.

- Katherine Campbell, L. A. (2003). The economic cost of publicly announced information security breaches: empirical evidence from the stock market. *Journal of Computer Security*(11), 431–448.
- Lawrence A. Gordon, M. P. (2011). The impact of information security breaches: Has there been a downward shift in costs? *Journal of Computer Security*(19), 33–56.
- Liu, W., Tanaka, H., & Matsuura, K. (n.d.). *An empirical analysis of security investment in countermeasures based on an enterprise survey in Japan (2006)*. State College, PA: Penn State University.
- Mandiant. (2013, February). *APT1—exposing one of China’s cyber units*. Retrieved June 6, 2014, from http://intelreport.mandiant.com/Mandiant_APT1_Report.pdf
- McAfee. (2011, February 10). *Global energy cyberattacks: “Night Dragon.”* Retrieved June 6, 2014, from <http://www.mcafee.com/us/resources/white-papers/wp-global-energy-cyberattacks-night-dragon.pdf>
- MITRE Corporation. (2012, October). *Cyber information-sharing models: an overview*. Retrieved June 2, 2014, from http://www.mitre.org/sites/default/files/pdf/cyber_info_sharing.pdf
- National Cyber League. (2013). *THE NATIONAL CYBER LEAGUE Fall 2013 Scoring*. (THE NATIONAL CYBER LEAGUE) Retrieved June 6, 2014, from <http://www.nationalcyberleague.org/scoring.shtml>
- Odlyzko, A. (2003). *Economics, Psychology, and Sociology of*. Minneapolis: University of Minnesota.
- Offensive Security Ltd. (2014). *Kali Linux*. (Offensive Security Ltd.) Retrieved June 6, 2014, from <http://www.kali.org/>
- O’Gorman, G., & McDonald, G. (2012, September 6). *The Elderwood project*. Retrieved June 6, 2014, from http://www.symantec.com/content/en/us/enterprise/media/security_response/whitepapers/the-elderwood-project.pdf
- Pal, R., Golubchik, L., & Psounis, K. (2011). Aegis - A novel cyber-insurance model. *Lecture Notes in Computer Science*, 7037, 131–150.
- Payne, S. C. (2006, June 19). *SANS Reading Room*. Retrieved June 6, 2014, from <http://www.sans.org/reading-room/whitepapers/auditing/guide-security-metrics-55>
- Pimendis, L. (2010, 6 11). *Cipher CTF—How To Play*. Retrieved June 24, 2014, from Cipher CTF: <http://www.cipher-ctf.org/HowToPlay.php>

- Project, T. H. (2014, June). *Honeynet project*. (The Honeynet Project) Retrieved June 6 , 2014, from <https://www.honeynet.org/>
- Rapid7. (2012, May 31). *Metasploitable 2 Exploitability Guide*. (Rapid7) Retrieved June 6, 2014, from <https://community.rapid7.com/docs/DOC-1875>
- Rogers, M. (2011). *Opening Statement for Open Hearing: Cyber threats and ongoing efforts to protect the nation*. Washington, D.C.
- Rogin, J. (2012, July 9). *The Cable*. Retrieved June 2, 2014, from http://thecable.foreignpolicy.com/posts/2012/07/09/nsa_chief_cybercrime_constitutes_the_greatest_transfer_of_wealth_in_history
- Scambray, J., McClure, S., & Kurtz, G. (2012). *Hacking exposed: network security secrets & solutions, 7th Edition*. McGraw-Hill Osborne Media.
- Security Onion. (2014). *Security Onion*. Retrieved June 6, 2014, from <http://blog.securityonion.net/>
- Symantec. (2014). *Advanced Persistent Threats: how they work*. (Symantec) Retrieved June 6, 2014, from <http://www.symantec.com/theme.jsp?themeid=apt-infographic-1>
- Symantec. (2014, January 20). *Symantec Cyber Readiness Challenge Player's Manual*. Retrieved June 6, 2014, from https://www.symantec.com/content/en/us/enterprise/other_resources/b-cyber-readiness-challenge-players-manual.pdf
- UnSpam Technologies. (2014). *Project Honeypot*. (UnSpam Technologies) Retrieved June 6, 2014, from <http://www.projecthoneypot.org/>
- Varian, H. (2000, June 1). Managing Online Security Risks. *New York Times*, p. 1.
- Verizon. (2014, June 6). *Verizon Enterprise Risk and Incident Sharing Metrics Framework*. Retrieved June 6, 2014, from http://www.verizonenterprise.com/resources/whitepapers/wp_verizon-incident-sharing-metrics-framework_en_xg.pdf
- Willemson, J. (2006). On the Gordon & Loeb Model for Information Security Investment. *Fifth Workshop on the Economics of Information Security, 2006*. Cambridge, UK.
- Yadron, D. (2014, May 4). Symantec Develops New Attack on Cyberhacking. *Wall Street Journal*.

THIS PAGE INTENTIONALLY LEFT BLANK

INITIAL DISTRIBUTION LIST

1. Defense Technical Information Center
Ft. Belvoir, Virginia
2. Dudley Knox Library
Naval Postgraduate School
Monterey, California